

UNITED STATES PATENT APPLICATION

for

METHOD AND SYSTEM FOR A FRONT-END MODULAR TRANSMISSION
CONTROL PROTOCOL (TCP) HANDOFF DESIGN IN A STREAMS BASED
TRANSMISSION CONTROL PROTOCOL INTERNET PROTOCOL (TCP/IP)
IMPLEMENTATION

Inventors:

Wenting Tang

Ludmila Cherkasova

Lance Russell

Prepared by:

WAGNER, MURABITO & HAO LLP

TWO NORTH MARKET STREET

THIRD FLOOR

SAN JOSE, CALIFORNIA 95113

(408) 938-9060

METHOD AND SYSTEM FOR A FRONT-END MODULAR TRANSMISSION
CONTROL PROTOCOL (TCP) HANDOFF DESIGN IN A STREAMS BASED
TRANSMISSION CONTROL PROTOCOL/INTERNET PROTOCOL (TCP/IP)
IMPLEMENTATION

5

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

The present invention relates to the field of STREAMS-
based Transmission Control Protocol/Internet Protocols
10 (TCP/IP) protocols. Specifically, the present invention
relates to the field of modular implementation of a TCP
handoff protocol in order to facilitate the transfer or
migration of TCP states from one node to another node in a
communication network. The present invention further
15 relates to the field of content-aware request distribution
in a web server cluster.

RELATED ART

Web server clusters are the most popular
20 configurations used to meet the growing traffic demands
imposed by the Internet. However, for web server clusters
to be able to achieve scalable performance, when the
cluster size increases, it is imperative that the cluster
employs some mechanism and/or policy for balanced request
25 distribution. For instance, it is important to protect the
web server clusters from overload and to provide service
differentiation when different client requests compete for
limited server resources. Mechanisms for intelligent
request distribution and request differentiation help to
30 achieve scalable and predictable cluster performance and

functionality, which are essential for today's Internet web sites.

Traditional request distribution methods try to
5 distribute the requests among the nodes in a web cluster
based on certain parameters, such as, IP addresses, port
numbers, and network load information. Some of these
request distribution methods have the ability to check the
packet header up to Layer 4 in the International
10 Organization for Standardization Open Systems
Interconnection (ISO/OSI) network reference model (e.g.,
TCP/IP) in order to make the distribution decision. As
such, these methods are commonly referred to as Layer 4
request distributions.

15
Figure 1 shows a communication network 100 of the prior
art that illustrates a load balancing solution. In Figure
1, a web server cluster 150 is shown. The cluster 150 can
be a web site with a virtual IP address located at the load
20 balancer 152. Various back-end web servers, such as back-
end web server-1 155, back-end web server-2 157, on up to
back-end web server-n 159 contain the content provided by
the web site.

25 Typically, the load-balancer 152 sits as a front-end
node on a local network and acts as a gateway for incoming
connections. The load balancer 152 is also called a request
distributor 152. Requests for content can come through the
Internet 120 from various clients, such as client-1 110,

client-2 112, on up to client-n 114. Incoming client requests are distributed, more or less, evenly to the pool of back-end web servers, without regard to the requested content. Further, the load balancer 152 forwards client requests to selected back-end nodes prior to establishing the connection with the client.

In this configuration, the web server cluster 150 appears as a single host to the clients. To the back-end web servers in a web cluster 150, the front-end load-balancer 152 appears as a gateway. In essence, it intercepts the incoming connection establishment packets and determines which back-end web server should process a particular request. Proprietary algorithms implemented in the front-end load balancer 152 are used to distribute the requests. These algorithms can take into account the number of back-end web servers available, the resources (CPU speed and memory) of each back-end web server, how many active TCP sessions are being serviced, etc. The balancing methods across different load-balancing servers vary, but in general, requests are forwarded to the least loaded back-end web server in the cluster 150.

In addition, only the virtual address located at the load balancer 152 is advertised to the Internet community, so the load balancer also acts as a safety net. The IP addresses of the individual back-end web servers are never sent back to the web browser located at the client making a request, such as client 110. The load-balancer rewrites

the virtual cluster IP address to a particular web server IP address using Network Address Translation (NAT).

However, because of this IP address rewriting, both
5 inbound requests and outbound responses must pass through the load-balancer 152. This creates a bottleneck and limits the scalability of the system.

Another solution for request distribution is
10 illustrated by the Brendel et al. patent (U.S. 5,774,660) by Resonate, Inc. In Brendel et al., a load balancer examines the content of the web request to provide for better efficiency in processing requests. However, the Brendel et al. patent platform weaves a proprietary protocol within the
15 TCP/IP protocol of an operating system of the load balancer. As a result, the algorithm utilized by the Brendel et al. patent necessitate kernel source modifications when porting from one operating system to another.

Also, in the Brendel et al. patent the proprietary
20 protocol is applied at the application layer of the operating system of the load balancer. Incoming packets to the load balancer have their protocol changed from TCP to a non-TCP (IXP) standard that is only understood by the
25 proprietary protocol located at the application layer. Later, the packets have their packets changed back to the TCP protocol for transmission to the back-end servers. Thus, the Brendel et al. patent reduces processing efficiency by switching back and forth between user level

00000632-061201

09880632.061201
kernels.. Further, were the Brendel et al. patent to be
implemented at the operating system's kernel level, any
modifications made to the proprietary protocol would
necessarily require access to the kernel source file which
5 typically is not available.

Thus, a need exists for more flexibility in designing
and implementing a TCP/IP handoff mechanism in a web server
cluster. Another need exists for a TCP/IP handoff mechanism
10 that is more portable between different operating systems
implementing the TCP/IP protocol. Still another need exists
for better efficiency in performing TCP/IP handoff
mechanisms.

SUMMARY OF THE INVENTION

Accordingly, a method and system for a method and system for a front-end modular Transmission Control Protocol (TCP) handoff design in a STREAMS-based Transmission Control Protocol/Internet Protocol (TCP/IP) implementation is described. Embodiments of the present invention provide for better management flexibility as TCP handoff (STREAMS) modules can be dynamically loaded and unloaded as dynamically loadable kernel modules (DLKM) without service interruption. In addition, embodiments of the present invention meet the above need as well as providing better portability between different operating systems since the TCP handoff modules can be ported to other STREAMS-enabled operating systems implementing the TCP/IP protocol. Also, embodiments of the present invention meet the above needs as well as providing for upper layer transparency in that no application modifications are necessary to take advantage of new solutions: modification are made at the kernel level in the DLKM TCP handoff modules possibly without shutting down the operating system. Further, embodiments of the present invention meet the above needs as well as providing for better efficiency in processing HTTP requests since the handoff modules only peek into message traffic with minimum functional replication of the original TCP/IP modules.

25

These and other objects and advantages of the present invention will no doubt become obvious to those of ordinary skill in the art after having read the following detailed

description of the preferred embodiments which are illustrated in the various drawing figures.

Specifically, the present invention discloses a method
5 and system for TCP handoff in a web cluster that includes a front-end node coupled to a plurality of back-end web servers. The web cluster can be a web site where the front-end node provides a virtual IP address for the web cluster. The content might be partitioned or replicated
10 between each of the back-end web servers. The back-end web servers may be coupled to the front-end node by a local area network. The web cluster is coupled to a communication network, such as the Internet.

15 The process begins by establishing a connection between a client web browser and the front-end node. The Bottom TCP (BTCP) handoff (STREAMS) module located below a TCP module in the operating system at the front-end node makes the connection with the client. The front-end BTCP
20 handoff module stores the connection setup messages for later TCP state migration purposes. The TCP module at the front-end node does not participate in making the connection or handing off of the TCP state.

25 The connection establishes a communication session between the client and the BTCP handoff module at the front-end node for the transfer of data contained within the content at the web site.

After the connection is made, the client sends a HTTP request to the BTCP handoff module at the front-end node. The BTCP handoff module parses the HTTP request by examining the content of the HTTP request to determine which of the back end servers is capable of or is most appropriate for processing the request. For example, the selected back-end web server contains the resources required by the HTTP request.

At this point, TCP state migration is made from the BTCP handoff module at the front-end node to the selected back-end web server containing the requested resources. The TCP state migration process occurs transparently to the client.

After successful migration of the TCP state to the selected back-end web server, a bottom IP (BIP) module in the front-end node switches to a forwarding mode, where packets received from the client are quickly forwarded to the selected back-end web server.

The purpose of the TCP state migration is to allow the selected back-end web server to act like the front-end node, actually the BTCP module at the front-end, for purposes of the communication session. The BTCP module at the selected back-end web server is able to change headers in outgoing packets from the back-end web server such that it appears that the source of the packets is the front-end node in its proper TCP state for the communication session.

After completely processing the HTTP request, the selected back-end web server terminates the communication session. In order to free the connection at all the nodes, 5 the BTCP handoff module at the selected back-end intercepts the TCP control packets, including the RST and FIN packets, and sends them to the BTCP module at the front-end node. The BTCP module at the front-end then sends the TCP control packets to the client. In this way, the connection is 10 terminated at all the nodes, including the client node, the front-end node, and the selected back-end web server.

In another embodiment of the present invention, the TCP state migration process includes storing the TCP/IP SYN 15 packet sent from said client to said front-end node to request a connection. The response TCP/IP SYN/ACK packet is sent from the front-end BTCP handoff module back to the client. An ACK packet from the client is also stored.

20 In order to handoff the TCP state to the selected back-end web server, the BTCP module at the front-end node replays the original TCP/IP SYN packet at the selected back-end web server. To the selected back-end web server it appears that a connection is being requested by a 25 client.

Contained within the TCP/IP SYN packet is the initial TCP sequence number of the BTCP module at the front-end node. The handoff protocol used by the handoff modules

allows the BTCP handoff module at the back-end web server
to read the information and update outgoing packets
appropriately, so that the source of the outgoing packets
appears to be the front-end node in the proper TCP state
5 for the communication session.

To complete the extension of the connection, the back-
end web server sends a TCP/IP SYN/ACK packet back to the
front-end node. The BTCP module at the selected back-end
10 updates the TCP/IP SYN/ACK packet to reflect the initial TCP
sequence number of the BTCP module at the front-end node.
An ACK packet is sent from the BTCP module at the front-end
node back to the selected back-end web server in order to
extend the connection. This ACK packet is identical to the
15 original ACK packet sent from the client except for changing
the destination address to the selected back-end web server.

BRIEF DESCRIPTION OF THE DRAWINGS

PRIOR ART Figure 1 illustrates a block diagram of an exemplary communication network implementing traditional load balancing solutions.

5

Figure 2 illustrates a block diagram of an exemplary communication network environment including a front-end web server coupled to a plurality of back-end web servers and the Internet for implementing a modular Transmission Control Protocol (TCP) handoff design in a STREAMS-based Transmission Control Protocol internet protocol (TCP/IP) implementation, in accordance with one embodiment of the present invention.

10

15

Figure 3A illustrates a block diagram of an exemplary STREAM-based modular framework for TCP/IP implementation, in accordance with one embodiment of the present invention.

20

Figure 3B illustrates a block diagram of the standard STREAMS-based modules used for TCP/IP implementation, in accordance with one embodiment of the present invention.

25

Figure 3C illustrates a block diagram of new STREAMS-based plug-in modules used for TCP handoff in STREAMS-based TCP/IP implementation, in accordance with one embodiment of the present invention.

Figure 4 illustrates a block diagram of an exemplary web server network environment including a front end server

coupled to a plurality of back-end web servers through a local area network, in accordance with one embodiment of the present invention.

5 Figure 5 illustrates a block diagram of an exemplary front-end network architecture that shows the request processing flow during a TCP handoff procedure, in accordance with one embodiment of the present invention.

10 Figure 6 is a flow diagram illustrating steps in a method for migrating TCP states from the front-end node to a selected back-end web server, in accordance with one embodiment of the present invention.

15 Figure 7 is a flow diagram illustrating steps in a method for establishing a connection between a client and a front-end node, in accordance with one embodiment of the present invention.

20 Figure 8 is a flow diagram illustrating steps in a method for extending a connection from a front-end node to a selected back-end web server, in accordance with one embodiment of the present invention.

09880632 .061201

DETAILED DESCRIPTION OF THE INVENTION

Reference will now be made in detail to the preferred
embodiments of the present invention, a method and system
for implementing TCP handoff in a STREAMS-based TCP/IP
5 implementation, examples of which are illustrated in the
accompanying drawings. While the invention will be
described in conjunction with the preferred embodiments, it
will be understood that they are not intended to limit the
invention to these embodiments. On the contrary, the
10 invention is intended to cover alternatives, modifications
and equivalents, which may be included within the spirit and
scope of the invention as defined by the appended claims.

Furthermore, in the following detailed description of
15 the present invention, numerous specific details are set
forth in order to provide a thorough understanding of the
present invention. However, it will be recognized by one of
ordinary skill in the art that the present invention may be
practiced without these specific details.

20

A method and system for a method and system for a
front-end modular Transmission Control Protocol (TCP)
handoff design in a STREAMS-based Transmission Control
Protocol/Internet Protocol (TCP/IP) implementation is
25 described. Embodiments of the present invention provide for
better management flexibility as TCP handoff (STREAMS)
modules can be dynamically loaded and unloaded as
dynamically loadable kernel modules (DLKM) without service
interruption. In addition, embodiments of the present

invention meet the above need as well as providing better portability between different operating systems since the TCP handoff modules can be ported to other operating systems with STREAMS-based TCP/IP implementation. Also, embodiments of the present invention meet the above needs as well as providing for upper layer transparency in that no application modifications are necessary to take advantage of new solutions: modification are made at the kernel level in the DLKM TCP handoff modules without accessing operating system source files. Further, embodiments of the present invention meet the above needs as well as providing for better efficiency in processing HTTP requests since the handoff modules only peek into message traffic with minimum functional replication of the original TCP/IP modules.

CONTENT AWARE REQUEST DISTRIBUTION

Content-aware request distribution takes into account the content (URL name, URL type, or cookies, etc.) when making a decision as to which back-end web server can best process the HTTP request. Content-aware request distribution mechanisms enable smart, specially tailored routing inside the web cluster.

Some benefits achieved in content-aware request distribution include allowing only partial replication of the content for a web site. Most, if not all, of the content provided by a web site server cluster can be partitioned. Additionally, the web site can further partition content based on specialization of information.

For example, dedicated back-end web servers can be set up to deliver different types of documents. Another benefit provided by content-aware distribution includes support for differentiated Web Quality of Service (Web QoS).

5

Content-aware request distribution can significantly improve the web cluster throughput. Content-aware request distribution based on cache affinity lead to significant performance improvements compared to the strategies taking
10 into account only load information.

Three main components comprise a web server cluster configuration in implementing a content-aware request distribution strategy: a dispatcher, a distributor, and a
15 web server. The dispatcher implements the request distribution strategy and decides which web server will be processing a given request. The distributor interfaces with the client and implements the TCP handoff in order to distribute the client requests to a specific web server.
20 The web server processes the client web requests, otherwise known as HTTP requests.

In the internet environment, the hypertext transfer protocol (HTTP) protocol is based on the connection-
25 oriented TCP protocol. In order to serve a client request (HTTP request packet), a TCP connection is first established between a client and the front-end node. A dispatcher component is accessed by the front-end node to determine which web server can process the HTTP request.

09880632.061201
The dispatcher component may be located at the front-end node or can be located at on a separate node.

In the present invention, the front-end node does not
5 serve the HTTP request. Therefore, some mechanism is
needed to forward the HTTP request for processing to the
right node in the web cluster. Embodiments of the present
invention utilize the STREAMS-based front-end modular TCP
handoff mechanism to support content-aware request
10 distribution.

The STREAMS-based front-end modular TCP handoff
mechanism enables the forwarding of back-end web server
responses directly to the clients without passing through
15 the front-end. Figure 2 illustrates an exemplary network
200 implementing the content-aware request distribution
utilizing a front-end server approach for TCP handoff, in
accordance with one embodiment of the present invention.
Figure 2 also is illustrative of the traffic flow between
20 the clients through the internet 220 to the web site
serviced by the web cluster 250.

Network 200 includes client-1 210, client-2 212, on up
to client-n 214. The clients are coupled to the web
25 cluster 250 via the internet 220. The web cluster includes
a front-end node 252 that does not provide any web content,
but distributes HTTP requests to a plurality of back-end
web servers, such as, back-end web server-1 255, back-end
web server-2 257, on up to back-end web server 259.

After the front-end node 252 establishes the connection with a client (e.g., client-1 210) and the request distribution decision is made, the established connection is
5 handed off to the selected back-end web server to serve the request. The TCP state, related to the established connection, is migrated from the front end to the selected back-end web server. One benefit of TCP handoff mechanism is that the back-end web servers can send response packets
10 directly to the client without routing outgoing packets back through the front-end node 252.

STREAMS-BASED TCP/IP IMPLEMENTATION

STREAMS-based TCP/IP implementation offers a framework
15 to implement the TCP handoff mechanism as plug-in modules in the TCP/IP stack, in accordance with one embodiment of the present invention.

The STREAMS modules provide the advantage of better
20 portability. The STREAMS-based modules are relatively independent of the original internal TCP/IP modules. In other words, STREAMS-based TCP handoff modules do not change any data structures or field values maintained by the original TCP/IP modules. Further, all the interactions
25 between TCP handoff modules and the original TCP/IP modules are messaged based, such that, no direct function calls are made. This enables maximum portability, so that designed TCP handoff modules can be ported to other STREAMS-based TCP/IP implementations very quickly.

Another advantage provided by the STREAMS-based modules is increased flexibility within the operating system. The TCP handoff modules may be dynamically loaded and unloaded as dynamically loadable kernel modules (DLKM) without service interruption. Improvements to the handoff mechanism are easily inserted as new TCP handoff modules into the kernel of an operating system without updating the operating system.

10

Furthermore, the STREAMS-based modules provide for increased efficiency when processing HTTP requests, especially in handing off TCP states from one node to another. The TCP handoff modules only peek into the TCP/IP message traffic. There is minimum functional replication of the original TCP/IP modules.

15

Also, the STREAMS-based modules allow for application transparency in the operating system. The TCP handoff mechanism operates at the kernel level within an operating system without any application layer involvement. Thus, no modifications at the application layer is necessary to perform TCP handoff. This is a valuable feature for applications where no source code is available.

20

25

Figure 3A illustrates a block diagram of a STREAMS-based modular framework for developing the communication services and the TCP handoff mechanism. Each stream generally has a stream head 310, a driver 314, and multiple

optional modules 312 between the stream head 310 and the driver 314. These modules 312 exchange information through messages. Messages can flow in the upstream direction or the downstream direction.

5

Each module 312 has a pair of queues: write queue and read queue. When a message passes through a queue, the routine for this queue is called to process the message. The routine can drop a message, pass a message, change the message header, or generate a new message.

10

The stream head 310 is responsible for interacting with the user processes 315. The stream head 310 accepts requests from the user processes 315, translates them into appropriate messages, and sends the messages downstream. The stream head 310 is also responsible for signaling to the user processes module 315 when new data arrives or some unexpected event happens.

15

Figure 3B illustrates a block diagram of the standard STREAMS-based modules used for TCP/IP STREAMS-based implementation, in accordance with one embodiment of the present invention. A transport provider interface (TPI) specification defines the message interface between the TCP module 320 and the stream head module 310. A data link provider interface (DLPI) specification defines the message interface between driver module 314 and the IP module 330. These two specifications, TPI and DLPI, can be implemented in individual STREAMS modules and define the message

20

25

format, valid sequences of messages, and semantics of messages exchanged between these neighboring modules.

For example, when the TCP module 320 receives SYN
5 request for establishing the HTTP connection, TCP module 320 sends a "T_CONN_IND" message upstream. Under the TPI specification, the TCP module 320 should not proceed until it gets the response from the application layer. However, in order to be compatible with BSD implementation-based
10 applications, TCP module 320 continues the connection establishment procedure with the client. When the application decides to accept the connection, it sends the "T_CONN_RES" downstream on the listen stream. It also creates another stream to accept this new connection, and
15 the TCP module 320 attaches a TCP connection state to this new stream. Data exchange continues on the accepted stream until either end closes the connection.

WEB SITE CLUSTER DESIGN FOR A FRONT-END ARCHITECTURE

20 As discussed previously, three main components comprise a web server cluster configuration in implementing a content-aware request distribution strategy: a dispatcher, a distributor, and a web server. The dispatcher implements the request distribution strategy and
25 decides which web server will be processing a given request. The distributor interfaces with the client and implements the TCP handoff in order to distribute the client requests to a specific web server. The web server processes the client requests, or HTTP requests.

09880632 061201

The HTTP protocol is stateless, in that each request is processed by the web server independently from the previous or subsequent request. However, in an e-commerce environment, the concept of a communication session defines a sequence of requests issued by the client and the responses to the request. For a communication session, it is important to maintain state information from previous interactions between a client and a server.

10

In one instance, state information allows for proper sequencing of data packets once they are received. In another instance, such state information might contain the content of the shopping cart or list of results from the search request. Thus, when the user is ready to make a purchase, or is asking for the next ten entries from the search list, the state information from the previous request should be retrieved. For efficient request processing and session integrity, it is desirable to send client requests to the same back-end web server. One of the popular schemes proposed for handling the state over the internet are cookies. Content-aware request routing provides a convenient mechanism to support session integrity, also known as a "sticky" session).

25

Figure 4 shows a web server cluster network 400 to support an exemplary e-commerce site implementing a front-end modular TCP handoff design, in accordance with one embodiment of the present invention. In this architecture,

the front-end node 410 has co-located the distributor
functionality 430 and the dispatcher functionality 420 used
to support the session integrity. For example, based on
the cookie attached to an incoming HTTP session, the front-
5 end node sends the HTTP request belonging to a particular
communication session to the initially chosen, same back-
end web server.

10 The specifics of the web cluster 400 illustrated in
Figure 4 is that the front-end node 410 and the back-end
web server nodes (server-1 442, server-2 444, on up to
server-n 446) have different functionalities. The front-
end node 410 combines the functions of distributor and
dispatcher. On the other hand, the back-end nodes perform
15 as the web servers that process the HTTP requests and
provide the web content.

The front-end node checks the cookie attached to an
HTTP request for a particular communication session, and
20 decides which back-end web server should process the
request. The distributor module 430 in the front end node
410 always hands off the connection to the appropriate
back-end web server, as the front-end node 410 never
processes the request. As such, this network architecture
25 400 is defined as by the term "TCP always-handoff."

Thus, the content-aware front-end node 410 in the
cluster network 400 shown in Figure 4 maintains the
communication session integrity mechanism. The front-end

node 410 accepts the connection, checks the HTTP header, and hands-off the connection to the proper back-end web server. The front-end node does not service the HTTP request.

5

MODULAR FRONT-END TCP HANDOFF DESIGN ARCHITECTURE

The TCP always-TCP handoff design is optimized for always TCP handoff of TCP states. As such, the TCP always-handoff design is optimized to minimize the TCP handoff procedures for remote processing of HTTP requests.

10

In always-TCP handoff, there are two kinds of nodes, the dedicated front-end node and the back-end web servers. The purpose of the front-end node is to trigger the client to send the HTTP request, and then handoff the connection to the selected server. Also, the TCP handoff mechanism, as shown in Figure 2, enables the forwarding of outgoing responses from the back-end web server nodes directly to the clients. Consequently, the front-end node is optimized to handoff the connection and forward the data as quickly as possible. The functionalities and the roles of the front-end and back-end web server nodes are different in this network architecture 200. As such, the set of STREAMS modules located the front-end and the back-end web servers provide different functionalities.

15

20

25

Two new modules are shown in Figure 3C to implement the functionality of TCP handoff at the front-end node, in accordance with one embodiment of the present invention.

According to the relative position in the existing TCP/IP stack located within the operating system of the front-end node, a bottom TCP (BTCP_{FE}) module 350 and a bottom IP (BIP_{FE}) module 360 are introduced. The BTCP_{FE} module 350 is located under the TCP_{FE} module 320 in the TCP/IP protocol stack. The BIP_{FE} module 360 is located under the IP_{FE} module 330 in the TCP/IP protocol stack. The TPI module 370 provides for the message interface between the TCP module 320 and the stream head module 310.

Each of the back-end web servers have only one additional STREAMS BTCP_{BE} module 520 that is located under a TCP_{BE} module 530, as is shown in Figure 5, in accordance with one embodiment of the present invention. An IP_{BE} module 510 is also located in the TCP/IP protocol stack implemented at each of the back end web servers.

A network connection 590 provides further communication between nodes in the web cluster including the front-end node and the back-end web servers as described in Figure 5. The network connection 590 can be over a local area network (LAN), a wide area network (WAN), or any suitable communication network including the Internet.

Figure 6 is a flow chart 600 illustrating steps in a method for migrating TCP states from the front-end node to a selected back-end web server, in accordance with one embodiment of the present invention.

Flow chart 600 begins with the front-end node finishing the three-way TCP handshake protocol used for connection establishment, in step 610. A client web browser initiates a connection request with the front-end node to establish a communication session for the transfer of data contained within the content available at the web site serviced by the front-end node. The front-end node provides a virtual IP address for a web cluster that includes a plurality of back-end web servers that contain the content provided by the web site. The content may be partitioned between each of the back-end web servers with little or no duplication. The front-end node also receives the requested document in step 610.

In step 620 of flow chart 600, the front-end node makes the routing decision. After the HTTP request is received at the front-end node, the HTTP request is parsed for content. A back-end web server is selected by the front-end node based on the content of the HTTP request. The back-end web server is selected to process the HTTP request.

In step 630 of flow chart 600, the front-end node migrates the TCP state of the front-end node from the front-end node to the selected back-end web server. The connection from the client to the web cluster must be extended from the front-end node to the selected back-end web server. This is accomplished by migrating the initial

state of the front-end node to the selected back-end. This is done through the TCP always-handoff modules at both the front-end and the back-end web servers. This process is transparent to the client. As such, the connection for the communication session extends from the client to the selected back-end web server.

Additionally, after successful TCP state migration, the front-end node enters a forwarding mode, in step 640. The front-end node forwards incoming data packets from the client to the selected back-end web server as quickly as possible.

In step 650 of flow chart 600, the front-end node terminates the forwarding mode and releases the related resources on the front-end node after the connection is closed.

Figure 7 is a flow chart 700 illustrating steps in a method for establishing a connection setup between a client and a front-end node, in accordance with one embodiment of the present invention. The BTCP_{FE} implements the connection setup function at the front-end node. Before the HTTP request is sent by the client to make a routing decision, the connection has to be established between the client and the front-end node. Packets are exchanged to establish the connection. After the connection is established, an HTTP request is passed from the client to the front-end node.

Since internet traffic follows a TCP/IP communication protocol, a TCP/IP SYN packet is sent from the client to the front-end node. The front-end node provides the
5 virtual IP address for a web site that includes a web cluster architecture including a front-end node and a plurality of back-end web servers, such as the network architectures shown in Figures 2 and 4.

10 The TCP/IP SYN packet arrives at the BTCP_{FE} module in step 710 of flow chart 700. At this stage, the BTCP_{FE} allocates a connection structure corresponding to each connection request upon receiving the TCP/IP SYN packet from the client.

15 In step 720 of flow chart 700, the BTCP_{FE} selects an initial sequence number according to its preference. In step 730, the BTCP_{FE} sends a TCP/IP SYN/ACK packet with the initial sequence number that, among other variables,
20 indicates the initial TCP state of the BTCP_{FE} module.

In step 740 of flow chart 700, the BTCP_{FE} module receives an TCP/IP ACK packet from the client. It is at this point that the connection is established between the
25 client and the front-end node. During this process, the BTCP_{FE} module emulates the TCP state transition and changes its TCP state accordingly.

09880632 061201
T02T90 2E908880

In step 750 of flow chart 700, the BTCP_{FE} module stores the connection packets (e.g., the SYN, SYN/ACK, and the ACK packets). These packets are replayed at a selected back-end web server to extend the connection for the communication session to the selected back-end web server.

In step 760 of flow chart 700, after the connection is established, the client sends the HTTP request to the front-end node. The HTTP request packets are intercepted and parsed by the BTCP_{FE} module and a back-end web server is selected based on the content of the request. It is important to note that all of the above activities in flow chart 700 occur without the involvement of the TCP module in the original operating system at the front-end node.

Once a back-end web server is selected to service the HTTP request, the connection for the communication session established by the HTTP request must be extended to the selected back-end web server. It is difficult to replicate this TCP state at the TCP module of the back-end web server directly. The TPI specification does not support schemes by which a new half-open TCP connection with a predefined state can be opened.

On the other hand, one embodiment of the present invention creates the half-open TCP connection by replaying the original connection packets to the TCP module (TCP_{BE}) at the selected back-end web server by the BTCP_{FE}. In

essence, the BTCP_{FE} acts as the client in some sense (see Figure 5).

Figure 8, is a flow chart 800 that in conjunction with Figure 5 illustrate steps in a method for extending the connection setup to a selected back-end web server, in accordance with one embodiment of the present invention.

In step 810 of flow chart 800, the first replayed TCP/IP SYN packet initiates the handoff process (see also step 2, Figure 5). The BTCP_{FE} module uses the connection establishment packets stored during the connection setup phrase and sends the TCP/IP SYN packet to the selected back-end web server. The destination address is changed in the TCP/IP SYN packet to reflect the selected back-end web server.

This TCP state migration occurs transparently to the client. Additionally, the connection is extended to the selected back-end web server. In other words, no separate connection is made between the front-end node and the selected back-end web server. To the selected back-end web server, it appears as if the connection sequence originated from the client instead of the front-end node. It is only the handoff modules that understand the correct addresses for the handoff procedure.

The TCP state migration is done by replaying the packets between the front-end node and the selected back-

end web server. Since all Internet traffic is handoff traffic, the TCP/IP SYN packet arriving at the listening port of the selected back-end web server indicates a handoff request is initiated.

5

In step 820 of flow chart 800, an initial sequence number is included within the TCP/IP SYN packet. This initial sequence number, and any other state information that may be included, indicates the initial TCP state of the BTCP_{FE} module.

10

The TCP/IP SYN packet is intercepted by the BTCP_{BE} module in step 830 of flow chart 800. The handoff protocol used by the handoff modules allows the bottom TCP (BTCP_{BE}) module at the back-end web server to read the initial sequence number, understand the initial state of the BTCP_{FE} module, and to update outgoing packets from the selected back-end web server to reflect the proper TCP state at the front-end node in the communication session.

15

20

In step 840 of flow chart 800, the TCP/IP SYN packet is sent upstream (see also step 2, Figure 5) to the TCP (TCP_{BE}) module at the selected back-end web server. The TCP_{BE} responds with a TCP/IP SYN/ACK packet and sends this downstream, back to the originating source, the client.

25

In step 850, the TCP/IP SYN/ACK packet is intercepted by the BTCP_{BE} module (see also step 3, Figure 5). In step 860, the BTCP_{BE} module changes the destination address to

the front-end node, and updates the TCP sequence number of the TCP/IP SYN/ACK packet to reflect the proper state of the front-end node used for establishing and maintaining the communication session.

5

In step 870, the TCP/IP SYN/ACK packet is sent to the front-end node where the BTCP_{FE} module intercepts the packet. In step 880, The BTCP_{FE} module then sends the updated ACK packet, used for establishing the connection between the client and the front-end node, back to the selected back-end web server (see also step 4, Figure 5). The HTTP request is forwarded over the network as in step 6 of Figure 5.

15 In another embodiment of the present invention, the initial sequence number of the front-end node is provided for in the original TCP/IP SYN packet that is sent to the selected back-end web server when extending the connection.

20 After handoff is processed successfully, the BTCP_{FE} module notifies the bottom IP (BIP_{FE}) module at the front-end node to switch to a forwarding mode for this connection. The BIP_{FE} continues to forward any packets on this connection until the forward session is closed. Packet forwarding should be done as quickly as possible. In this configuration, it might be better to forward the packet on top of the device driver. Additionally, data forwarding may happen at Layer 2 or Layer 3, depending on the cluster configuration.

When the forwarding packets need to traverse a router or across a WAN, the packet's destination address may be changed to the selected back-end web server's IP address.

5 In return, packets coming from the selected back-end web server may update the source IP address to duplicate the front-end node's virtual IP address.

Further, the BIP_{FE} module updates the destination IP
10 address for incoming packets to the selected server's IP address, and sends the packet upstream. The IP_{FE} forwards the packet according to its routing tables to the selected back-end node. The BIP_{BE} module also intercepts the
15 incoming packets and manipulates the TCP header anyway to update the initial sequence number and the TCP checksum.

The handoff connection is closed in the following fashion. The BTCP_{BE} intercepts the TCP control packets (e.g., packets with RST or FIN flags) and sends it to the
20 BTCP_{FE} (see step 7, Figure 5). The BTCP_{FE} records the connection progress, updates, and relays the packets to the client. As usual, data traffic goes directly to the client. Additionally, the front-end sees two way traffic and may keep track of the handoff connection status and
25 close the connection in a timely manner.

While the methods of embodiments illustrated in flow charts 600, 700, and 800 show specific sequences and quantity of steps, the present invention is suitable to

alternative embodiments. For example, not all the steps provided for in the method are required for the present invention. Furthermore, additional steps can be added to the steps presented in the present embodiment. Likewise,
5 the sequences of steps can be modified depending upon the application.

Embodiments of the present invention, a method and system for a front-end modular Transmission Control Protocol
10 (TCP) handoff design in a STREAMS-based Transmission Control Protocol Internet protocol (TCP/IP) implementation, is thus described. While the present invention has been described in particular embodiments, it should be appreciated that the present invention should not be construed as limited by such
15 embodiments, but rather construed according to the below claims.